

Fast Algorithm for Partial Covers in Words*

Tomasz Kociumaka^{†1}, Jakub Radoszewski^{‡1}, Wojciech Rytter^{§1,2}, Solon P. Pissis^{¶3}, and
Tomasz Walen^{||1}

¹Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland,
[kociumaka, jrad, rytter, walen]@mimuw.edu.pl

²Faculty of Mathematics and Computer Science, Copernicus University, Toruń, Poland

³Department of Informatics, King's College London, UK, solon.pissis@h-its.org

Abstract

A factor u of a word w is a *cover* of w if every position in w lies within some occurrence of u in w . A word w covered by u thus generalizes the idea of a *repetition*, that is, a word composed of exact concatenations of u . In this article we introduce a new notion of α -*partial cover*, which can be viewed as a relaxed variant of cover, that is, a factor covering at least α positions in w . We develop a data structure of $\mathcal{O}(n)$ size (where $n = |w|$) that can be constructed in $\mathcal{O}(n \log n)$ time which we apply to compute all shortest α -partial covers for a given α . We also employ it for an $\mathcal{O}(n \log n)$ -time algorithm computing a shortest α -partial cover for each $\alpha = 1, 2, \dots, n$.

1 Introduction

The notion of periodicity in words and its many variants have been well-studied in numerous fields like combinatorics on words, pattern matching, data compression, automata theory, formal language theory, and molecular biology (see [9]). However the classic notion of periodicity is too restrictive to provide a description of a word such as `abaababaaba`, which is covered by copies of `aba`, yet not exactly periodic. To fill this gap, the idea of *quasiperiodicity* was introduced [1]. In a periodic word, the occurrences of the period do not overlap. In contrast, the occurrences of a quasiperiod in a quasiperiodic word may overlap. Quasiperiodicity thus enables the detection of repetitive structures that would be ignored by the classic characterization of periods.

The most well-known formalization of quasiperiodicity is the cover of word. A factor u of a word w is said to be a *cover* of w if $u \neq w$, and every position in w lies within some occurrence of u in w . Equivalently, we say that u *covers* w . Note that a cover of w must also be a *border* — both prefix and suffix — of w . Thus, in the above example, `aba` is the shortest cover of `abaababaaba`.

A linear-time algorithm for computing the shortest cover of a word was proposed by Apostolico et al. [2], and a linear-time algorithm for computing all the covers of a word was proposed by Moore & Smyth [23]. Breslauer [4] gave an online linear-time algorithm computing the *minimal cover array* of a word — a data structure specifying the shortest cover of every prefix of the word. Li & Smyth [22] provided a linear-time

*A preliminary version of this work appeared in the Proceedings of the Twenty-Fourth Annual Symposium on Combinatorial Pattern Matching, pp. 177–188, 2013.

[†]Supported by Polish budget funds for science in 2013–2017 as a research project under the ‘Diamond Grant’ program.

[‡]The author receives financial support of Foundation for Polish Science.

[§]Supported by grant no. N206 566740 of the National Science Centre.

[¶]Supported by the NSF-funded iPlant Collaborative (NSF grant #DBI-0735191).

^{||}Supported by Iuventus Plus grant (IP2011 058671) of the Polish Ministry of Science and Higher Education.

algorithm for computing the *maximal cover array* of a word, and showed that, analogous to the border array [8], it actually determines the structure of *all* the covers of every prefix of the word.

A known extension of the notion of cover is the notion of *seed*. A seed is not necessarily aligned with the ends of the word being covered, but is allowed to overflow on either side. More formally, a word u is a seed of w if u is a factor of w and w is a factor of some word y covered by u . Seeds were first introduced by Iliopoulos, Moore, and Park [18]. A linear algorithm for computing the shortest seed of a word was given by Kociumaka et al. [19].

Still it remains unlikely that an arbitrary word, even over the binary alphabet, has a cover (or even a seed). For example, `abaaababababaaababaa` is a word that not only has no cover, but whose every prefix also has no cover. In this article we provide a natural form of quasiperiodicity. We introduce the notion of *partial covers*, that is, factors covering at least a given number of positions in w . Recently, Flouri et al. [14] suggested a related notion of *enhanced covers* which are additionally required to be borders of the word.

Partial covers can be viewed as a relaxed variant of covers alternative to approximate covers [24]. The approximate covers require each position to lie within an approximate occurrence of the cover. This allows for small irregularities within each fragment of a word. On the other hand partial covers require exact occurrences but drop the condition that all positions need to be covered. This allows some fragments to be completely irregular as long as the total length of such fragments is small. The significant advantage of partial covers is that they enjoy more combinatorial properties, and consequently the algorithms solving the most natural problems are much more efficient than those concerning approximate covers, where the time complexity rarely drops below quadratic and some problems are even NP-hard.

Let $Covered(u, w)$ denote the number of positions in w covered by occurrences of the word u in w ; we call this value the *cover index* of u within w . For example, $Covered(\text{aba}, \text{aababab}) = 5$. We primarily focus on the following two problems, but the tools we develop can be used to answer a number of questions concerning partial covers, some of which are discussed in the Conclusions.

PARTIALCOVERS problem

Input: a word w of length n and a positive integer $\alpha \leq n$.

Output: all shortest factors u such that $Covered(u, w) \geq \alpha$.

Each factor given in the output is represented by the first and the last starting position of its occurrence in w .

Example 1. Let $w = \text{bccacccaccacccb}$ and $\alpha = 11$. Then the only shortest α -partial covers are `ccac` and `cacc`.

ALLPARTIALCOVERS problem

Input: a word w of length n .

Output: for all $\alpha = 1, \dots, n$, a shortest factor u such that $Covered(u, w) \geq \alpha$.

Our contribution. The following summarizes our main result.

Theorem 1. *The PARTIALCOVERS and ALLPARTIALCOVERS problems can be solved in $\mathcal{O}(n \log n)$ time and $\mathcal{O}(n)$ space.*

We extensively use suffix trees, for an exposition see [8, 12]. A suffix tree of a word is a compact trie of its suffixes, the nodes of the trie which become nodes of the suffix tree are called *explicit* nodes, while the other nodes are called *implicit*. Each edge of the suffix tree can be viewed as an *upward* maximal path of implicit nodes starting with an explicit node. Moreover, each node belongs to a unique path of that kind. Then, each node of the trie can be represented in the suffix tree by the edge it belongs to and an index within the corresponding path. Each factor of the word corresponds to an explicit or implicit node of the suffix tree. A representation of this node is called the *locus* of the factor. Our algorithm finds the loci of the shortest partial covers, it is then straightforward to locate an occurrence for each of them.

A Sketch of the Algorithm The algorithm first augments the suffix tree of w , that is, a linear number of implicit extra nodes become explicit. Then, each node of the augmented tree is annotated with two integer values. They allow for determining the size of the covered area for each implicit node by a simple formula, since limited to a single edge of the augmented suffix tree, these values form an arithmetic progression. This yields a solution to the PARTIALCOVERS. For an efficient solution to the ALLPARTIALCOVERS problem, we additionally find the upper envelope of a number of line segments constructed from the arithmetic progressions.

Structure of the Paper In Section 2 we formally introduce the augmented and annotated suffix tree that we call *Cover Suffix Tree*. We show its basic properties and present its application for PARTIALCOVERS and ALLPARTIALCOVERS problems. Section 4 is dedicated to the construction of the Cover Suffix Tree. Before that, Section 3 presents an auxiliary data structure being an extension of the classical Union/Find data structure; its implementation is given later, in Section 5. Additional applications of the Cover Suffix Tree are given in Sections 6 and 7. The former presents how the data structure can be used to compute all primitively rooted squares in a word and a linear-sized representation of all the seeds in a word. The latter contains a short discussion of variants of the PARTIALCOVERS problem that can be solved in a similar way.

2 Augmented and Annotated Suffix Trees

Let w be a word of length n over a totally ordered alphabet Σ . The suffix tree T of w can be constructed in $\mathcal{O}(n \log |\Sigma|)$ time [13, 25]. For an explicit or implicit node v of T , we denote by \hat{v} the word obtained by spelling the characters on a path from the root to v . We also denote $|v| = |\hat{v}|$. As in most applications of the suffix tree, the leaves of T play an auxiliary role and do not correspond to factors (actually they are suffixes of $w\#$, where $\# \notin \Sigma$). They are labeled with the starting positions of the suffixes of w .

We introduce the *Cover Suffix Tree* of w , denoted by $CST(w)$, as an *augmented* — new nodes are added — suffix tree in which the nodes are *annotated* with information relevant to covers. $CST(w)$ is similar to the data structure named *Minimal Augmented Suffix Tree* (see [3, 5]).

For a set X of integers and $x \in X$, we define

$$next_X(x) = \min\{y \in X, y > x\},$$

and we assume $next_X(x) = \infty$ if $x = \max X$. By $Occ(v, w)$ we denote the set of starting positions of occurrences of \hat{v} in w . For any $i \in Occ(v, w)$, we define:

$$\delta(i, v) = next_{Occ(v, w)}(i) - i.$$

Note that $\delta(i, v) = \infty$ if i is the last occurrence of \hat{v} . Additionally, we define:

$$cv(v) = Covered(\hat{v}, w), \quad \Delta(v) = |\{i \in Occ(v, w) : \delta(i, v) \geq |v|\}|;$$

see, for example, Fig. 1.




Figure 1: Let $w = \text{bccacccaccaccc}$ and let v be the node corresponding to cacc . We have $Occ(v, w) = \{4, 8, 11\}$, $cv(v) = 11$, $\Delta(v) = 2$.

A word u is called *primitive* if $u = y^k$ for a word y and an integer k implies that $y = u$, and non-primitive otherwise. A square u^2 is called *primitively rooted* if u is primitive.

Observation 1. Let v be a node in the suffix trie of w . Then $\hat{v}\hat{v}$ is a primitively rooted square in w if and only if there exists $i \in Occ(v, w)$ such that $\delta(i, v) = |v|$.

Proof. Recall that, by the synchronization property of primitive words (see [8]), \hat{v} is primitive if and only if it occurs exactly twice in $\hat{v}\hat{v}$.

(\Rightarrow) If $\hat{v}\hat{v}$ occurs in w at position i then $\delta(i, v) = |v|$.

(\Leftarrow) If $\delta(i, v) = |v|$ then obviously $\hat{v}\hat{v}$ occurs in w at position i . Additionally, if \hat{v} was not primitive then $\delta(i, v) < |v|$ would hold. \square

In $CST(w)$, we introduce additional explicit nodes called *extra nodes*, which correspond to halves of primitively rooted square factors of w . Moreover we annotate all explicit nodes (including extra nodes) with the values cv, Δ ; see, for example, Fig. 2. The number of extra nodes is bounded by the number of distinct squares, which is linear [15], so $CST(w)$ takes $\mathcal{O}(n)$ space.

Lemma 1. *Let v_1, v_2, \dots, v_k be the consecutive implicit nodes on the edge from an explicit node v of $CST(w)$ to its explicit parent. Then for $1 \leq i \leq k$ we have*

$$cv(v_i) = cv(v) - i\Delta(v),$$

in particular $(cv(v_i))_{i=1}^k$ forms an arithmetic progression.

Proof. Note that $Occ(v_i, w) = Occ(v, w)$, since otherwise v_i would be an explicit node of $CST(w)$. Also note that if any two occurrences of \hat{v} in w overlap, then the corresponding occurrences of \hat{v}_i overlap. Otherwise, by Observation 1, the path from v to v_i (excluding v) would contain an extra node. Hence, when we go up from v (before reaching its parent) the size of the covered area decreases at each step by $\Delta(v)$. \square




Figure 2: $CST(w)$ for $w = bcccacccaccacccb$. It contains four extra nodes that are denoted by squares in the figure. Each node is annotated with $cv(v), \Delta(v)$. Leaves are omitted for clarity.

Example 2. Consider the word w from Fig. 2. The word $cccacc$ corresponds to an explicit node of $CST(w)$; we denote it by v . We have $cv(v) = 10$ and $\Delta(v) = 1$ since the two occurrences of the factor $cccacc$ in w overlap. The word $cccac$ corresponds to an implicit node v' and $cv(v') = 10 - 1 = 9$. Now the word $ccca$ corresponds to an extra node v'' of $CST(w)$. Its occurrences are adjacent in w and $cv(v'') = 8, \Delta(v'') = 2$. The word ccc corresponds to an implicit node v''' and $cv(v''') = 8 - 2 = 6$.

As a consequence of Lemma 1 we obtain the following result. Recall that the locus of a factor v of w , given by its start and end position in w , can be found in $\mathcal{O}(\log \log |v|)$ time [21].

Lemma 2. *Assume we are given $CST(w)$. Then we can compute:*

- (1) *for any α , the loci of the shortest α -partial covers in linear time;*
- (2) *given the locus of a factor u in the suffix tree $CST(w)$, the cover index $Covered(u, w)$ in $\mathcal{O}(1)$ time.*

Proof. Part (2) is a direct consequence of Lemma 1. As for part (1), for each edge of $CST(w)$, leading from v to its parent v' , we need to find minimum $|v| \geq j > |v'|$ for which $cv(v) - \Delta(v) \cdot (|v| - j) \geq \alpha$. Such a linear inequality can be solved in constant time. \square

Due to this fact the efficiency of the PARTIALCOVERS problem relies on the complexity of $CST(w)$ construction. In turn, the following lemma, also a consequence of Lemma 1, can be used to solve ALLPARTIALCOVERS problem provided that $CST(w)$ is given. As a tool a solution to the geometric problem of upper envelope [17] is applied.

Lemma 3. *Assume we are given $CST(w)$. Then we can compute the locus of a shortest α -partial cover for each $\alpha = 1, 2, \dots, n$ in $\mathcal{O}(n \log n)$ time and $\mathcal{O}(n)$ space.*

Proof. Consider an edge of $CST(w)$ from v to its parent v' containing k implicit nodes. For each such edge, we form a line segment on the plane connecting points $(|v|, cv(v))$ and $(|v| - k, cv(v) - k \cdot \Delta(v))$ (if there are no implicit nodes on the edge, the line segment is a single point). Denote all such line segments obtained from $CST(w)$ as s_1, \dots, s_m , we have $m = \mathcal{O}(n)$. We consider the upper envelope \mathcal{E} of the set of these segments. Formally, if each s_i connecting points (x_i, y_i) and (x'_i, y'_i) , $x_i \leq x'_i$, is interpreted as a linear function on a domain $[x_i, x'_i]$, \mathcal{E} is defined as a function $\mathcal{E} : [1, n] \rightarrow [1, n]$ such that:

$$\mathcal{E}(x) = \max\{s_i(x) : i \in \{1, \dots, m\}, x \in [x_i, x'_i]\}.$$

Here we are actually interested in an *integer envelope* \mathcal{E}' , that is, \mathcal{E} limited to integer arguments, see Fig. 3. By Lemma 1, for any $j \in \{1, \dots, n\}$, $\mathcal{E}'(j)$ equals the maximum of $Covered(u, w)$ over all factors u of w such that $|u| = j$. A piecewise linear representation of \mathcal{E} can be computed in $\mathcal{O}(m \log m)$ time and $\mathcal{O}(m)$ space [17], therefore the function \mathcal{E}' for all its arguments can be computed in the same time complexity.

Let us introduce a prefix maxima sequence for \mathcal{E}' : $\mu_i = \max\{\mathcal{E}'(j) : j \in \{1, \dots, i\}\}$, with $\mu_0 = 0$. Note that μ_i is non-decreasing. If $\mu_i > \mu_{i-1}$ then the shortest α -partial cover for all $\alpha \in (\mu_{i-1}, \mu_i]$ has length i . An example of such a partial cover can be recovered if we explicitly store the initial line segments used in the pieces of the representation of \mathcal{E} . Thus the solution of the ALLPARTIALCOVERS problem can be obtained from the sequence μ_i in $\mathcal{O}(m) = \mathcal{O}(n)$ time. \square

In the following two sections we provide an $\mathcal{O}(n \log n)$ time construction of $CST(w)$. Together with Lemmas 2 and 3, it yields Theorem 1.

3 Extension of Disjoint-Set Data Structure

In this section we extend the classic disjoint-set data structure to compute the *change lists* of the sets being merged, as defined below. First, let us extend the *next* notation. For a partition $\mathcal{P} = \{P_1, \dots, P_k\}$ of $U = \{1, \dots, n\}$, we define

$$next_{\mathcal{P}}(x) = next_{P_i}(x) \text{ where } x \in P_i.$$

Now for two partitions $\mathcal{P}, \mathcal{P}'$ let us define the *change list* (see also Fig. 4) by

$$ChangeList(\mathcal{P}, \mathcal{P}') = \{(x, next_{\mathcal{P}'}(x)) : next_{\mathcal{P}}(x) \neq next_{\mathcal{P}'}(x)\}.$$

We say that (\mathcal{P}, id) is a partition of U labeled by L if \mathcal{P} is a partition of U and $id : \mathcal{P} \rightarrow L$ is a one-to-one (injective) mapping. A label $\ell \in L$ is called *active* if $id(P) = \ell$ for some $P \in \mathcal{P}$ and *free* otherwise.




Figure 3: Line segments constructed as in Lemma 3 for the $CST(w)$ from Fig. 2. The marked points joined with a dashed polyline show the values of the integer upper envelope function \mathcal{E}' . We infer from the graph that the lengths of the shortest α -partial covers of w are as follows: 1 for $\alpha \leq 10$, 4 for $\alpha = 11$, 5 for $\alpha = 12$, and α for $\alpha \geq 13$.




Figure 4: Let \mathcal{P} be the partition of $\{1, \dots, 9\}$ whose classes consist of leaves in the subtrees rooted at children of v , $\mathcal{P} = \{\{1, 3, 4\}, \{2, 5, 6, 7\}, \{8, 9\}\}$, and let $\mathcal{P}' = \{\{1, \dots, 9\}\}$. Then $ChangeList(\mathcal{P}, \mathcal{P}') = \{(1, 2), (2, 3), (4, 5), (7, 8)\}$ (depicted by dotted arrows).

Lemma 4. *Let $n \leq k$ be positive integers such that k is of magnitude $\Theta(n)$. There exists a data structure of size $\mathcal{O}(n)$, which maintains a partition (\mathcal{P}, id) of $\{1, \dots, n\}$ labeled by $L = \{1, \dots, k\}$ and supports the following operations:*

- *Find(x) for $x \in \{1, \dots, n\}$ gives the label of $P \in \mathcal{P}$ containing x .*
- *Union(I, ℓ) for a set I of active labels and a free label ℓ replaces all $P \in \mathcal{P}$ with labels in I by their set-theoretic union with the label ℓ . The change list of the corresponding modification of \mathcal{P} is returned.*

Initially \mathcal{P} is a partition into singletons with $id(\{x\}) = x$. Any valid sequence of Union operations is performed in $\mathcal{O}(n \log n)$ time. A single Find operation takes $\mathcal{O}(1)$ time.

Note that these are actually standard disjoint-set data structure operations except for the fact that we require *Union* to return the change list. The technical proof of Lemma 4 is postponed until Section 5.

4 $\mathcal{O}(n \log n)$ -time Construction of $CST(w)$

The suffix tree of w augmented with extra nodes is called the *skeleton* of $CST(w)$, which we denote by $sCST(w)$. It could be constructed using the fact that all square factors of a word can be computed in linear time [16, 10, 11]. However, we do not need such a complicated machinery here. We will compute $sCST(w)$ on the fly, simultaneously annotating the nodes with cv , Δ .

We introduce auxiliary notions related to covered area of nodes:

$$cv_h(v) = \sum_{\substack{i \in Occ(v, w) \\ \delta(i, v) < h}} \delta(i, v), \quad \Delta_h(v) = |\{i \in Occ(v, w) : \delta(i, v) \geq h\}|.$$

Observation 2. $cv(v) = cv_{|v|}(v) + \Delta_{|v|}(v) \cdot |v|$, $\Delta(v) = \Delta_{|v|}(v)$.

In the course of the algorithm some nodes will have their values c, Δ already computed; we call them *processed nodes*. Whenever v will be processed, so will its descendants.

The algorithm processes inner nodes v of $sCST(w)$ in the order of non-increasing height $h = |v|$. The height is not defined for leaves, so we start with $h = n + 1$. Extra nodes are created on the fly using Observation 1 (this takes place in the auxiliary *Lift* routine).

We maintain the partition \mathcal{P} of $\{1, \dots, n\}$ given by sets of leaves of subtrees rooted at *peak nodes*. Initially the peak nodes are the leaves of $sCST(w)$. Each time we process v all its children are peak nodes. Consequently, after processing v they are no longer peak nodes and v becomes a new peak node. The sets in the partition are labeled with identifiers of the corresponding peak nodes. Recall that leaves are labeled with the starting positions of the corresponding suffixes. We allow any labeling of the remaining nodes as long as each node of $sCST(w)$ has a distinct label of magnitude $\mathcal{O}(n)$. For this set of labels we store the data structure of Lemma 4 to compute the change list of the changing partition.




Figure 5: One stage of the algorithm, where the peak nodes are v_1, \dots, v_5 while the currently processed node is v . If $i \in List[d]$ and $v_3 = Find(i)$, then $d = \delta(i, v_3) = Dist[i]$. The current partition is $\mathcal{P} = \{Leaves(v_1), Leaves(v_2), Leaves(v_3), Leaves(v_4), Leaves(v_5)\}$. After v is processed, the partition changes to $\mathcal{P} = \{Leaves(v_1), Leaves(v_2), Leaves(v), Leaves(v_5)\}$. The *Union* operation merges $Leaves(v_4), Leaves(v_3)$ and returns the corresponding change list.

We maintain the following technical invariant (see Fig. 5).

Invariant(h):

(A) For each peak node z we store:

$$cv'[z] = cv_h(z), \quad \Delta'[z] = \Delta_h(z).$$

(B) For each $i \in \{1, \dots, n\}$ we store $Dist[i] = \delta(i, Find(i))$.

(C) For each $d < h$ we store $List[d] = \{i : Dist[i] = d\}$.

We use two auxiliary routines. The *Lift* operation updates cv' and Δ' values when h decrements. It also creates all extra nodes of depth h . The *LocalCorrect* operation is used for updating cv' and Δ' values for children of the node v . The *Dist* and *List* arrays are stored to enable efficient implementation of these two routines.

```

Algorithm COMPUTECST( $w$ )
   $T :=$  suffix tree of  $w$ ;
   $\mathcal{P} :=$  partition of  $\{1, \dots, n\}$  into singletons;
  foreach  $v : a \text{ leaf of } T$  do  $cv'[v] := 0, \Delta'[v] := 1$ ;
  for  $h := n + 1$  downto  $0$  do
    Lift( $h$ );
    {Now part (A) of Invariant( $h$ ) is satisfied}
    foreach  $v : an \text{ inner node of } T, |v| = h$  do
       $cv'[v] := \sum_{u \in children(v)} cv'[u]$ ;
       $\Delta'[v] := \sum_{u \in children(v)} \Delta'[u]$ ;
       $ChangeList(v) := Union(children(v), v)$ 
      foreach  $(p, q) \in ChangeList(v)$  do LocalCorrect( $p, q, v$ );
       $cv[v] := cv'[v] + \Delta'[v] \cdot |v|$ ;  $\Delta[v] := \Delta'[v]$ ;
  return  $T$  together with values of  $cv, \Delta$ ;

```

Description of the *Lift*(h) Operation The procedure *Lift* plays an important preparatory role in processing the current node. According to part (A) of our invariant, for all peak nodes z we know the values: $cv'[z] = cv_{h+1}(z)$, $\Delta'[z] = \Delta_{h+1}(z)$. Now we have to change $h+1$ to h and guarantee validity of the invariant: $cv'[z] = cv_h(z)$, $\Delta'[z] = \Delta_h(z)$. This is exactly how the following operation updates cv' and Δ' .

It also creates all extra nodes of depth h that were not explicit nodes of the suffix tree. By Observation 1, if $i \in List[h]$ then at position i in w there is an occurrence of a primitively rooted square of half length h . Consequently, an extra node corresponding to this occurrence is created in the *Lift* operation.

```

Function Lift( $h$ )
  foreach  $i$  in  $List[h]$  do
     $v := Find(i)$ ;
     $\Delta'[v] := \Delta'[v] + 1$ ;  $cv'[v] := cv'[v] - h$ ;
    if  $|parent(v)| < h$  then
      Create a node of depth  $h$  on the edge from
       $parent(v)$  to  $v$ 

```

Description of the *LocalCorrect*(p, q, v) Operation Here we assume that \hat{v} occurs at positions $p < q$ and that these are consecutive occurrences. Moreover, we assume that these occurrences are followed by distinct characters, i.e. $(p, q) \in ChangeList(v)$. The *LocalCorrect* procedure updates $Dist[p]$ to make part (B) of the invariant hold for p again. The data structure *List* is updated accordingly so that (C) remains satisfied.

Function *LocalCorrect*(p, q, v)

$d := q - p$; $d' := \text{Dist}[p]$;

if $d' < |v|$ **then** $cv'[v] := cv'[v] - d'$ **else** $\Delta'[v] := \Delta'[v] - 1$;

if $d < |v|$ **then** $cv'[v] := cv'[v] + d$ **else** $\Delta'[v] := \Delta'[v] + 1$;

$\text{Dist}[p] := d$;

$\text{remove}(i, \text{List}[d'])$; $\text{insert}(i, \text{List}[d])$;

Complexity of the Algorithm In the course of the algorithm we compute $\text{ChangeList}(v)$ for each $v \in T$. Due to Lemma 4 we have:

$$\sum_{v \in T} |\text{ChangeList}(v)| = \mathcal{O}(n \log n).$$

Consequently we perform $\mathcal{O}(n \log n)$ operations *LocalCorrect*. In each of them at most one element is added to a list $\text{List}[d]$ for some d . Hence the total number of insertions to these lists is also $\mathcal{O}(n \log n)$.

The cost of each operation *Lift* is proportional to the total size of the list $\text{List}[h]$ processed in this operation. For each h the list $\text{List}[h]$ is processed once and the total number of insertions into lists is $\mathcal{O}(n \log n)$, therefore the total cost of all operations *Lift* is also $\mathcal{O}(n \log n)$. This proves the following fact which, together with Lemmas 2 and 3, implies our main result (Theorem 1).

Lemma 5. *Algorithm COMPUTECST constructs $\text{CST}(w)$ in $\mathcal{O}(n \log n)$ time and $\mathcal{O}(n)$ space, where $n = |w|$.*

5 Implementation Details

In this section we give a proof of Lemma 4. We use an approach similar to Brodal and Pedersen [6] (who use the results of [7]) originally devised for computation of maximal quasiperiodicities.

Theorem 3 of [6] states that a subset X of a linearly ordered universe can be stored in a height-balanced tree of linear size supporting the following operations:

$X.\text{MultiInsert}(Y)$: insert all elements of Y to X ,

$X.\text{MultiPred}(Y)$: return all (y, x) for $y \in Y$ and $x = \max\{z \in X, z < y\}$,

$X.\text{MultiSucc}(Y)$: return all (y, x) for $y \in Y$ and $x = \min\{z \in X, z > y\}$,

in $\mathcal{O}\left(|Y| \max\left(1, \log \frac{|X|}{|Y|}\right)\right)$ time.

In the data structure we store each $P \in \mathcal{P}$ as a height-balanced tree. Additionally, we store several auxiliary arrays, whose semantics follows. For each $x \in \{1, \dots, n\}$ we maintain a value $\text{next}[x] = \text{next}_P(x)$ and a pointer $\text{tree}[x]$ to the tree representing P such that $x \in P$. For each $P \in \mathcal{P}$ (technically for each tree representing $P \in \mathcal{P}$) we store $\text{id}[P]$ and for each $\ell \in L$ we store $\text{id}^{-1}[\ell]$, a pointer to the corresponding tree (null for free labels).

Answering *Find* is trivial as it suffices to follow the *tree* pointer and return the *id* value. The *Union* operation is performed according to the pseudocode given below (for brevity we write P_i instead of $\text{id}^{-1}[i]$).

Claim 2. *The Union operation correctly computes the change list and updates the data structure.*

Proof. In the *Union* operation for sets $P_i, i \in I$, we find the largest set P_{i_0} and *MultiInsert* all the elements of the remaining sets to P_{i_0} . If (a, b) is in the change list, then a and b come from different sets P_i , in particular at least one of them does not come from P_{i_0} . Depending on which one it is, the pair (a, b) is found by *MultiPred* or *MultiSucc* operation. While computing C , the table *next* is not updated yet (i.e. corresponds to the state before *Union* operation) while S is already updated. Consequently the pairs inserted to C indeed belong to the change list. Once C is proved to be the change list, it is clear that *next* is updated correctly. For the other components of the data structure, correctness of updates is evident. □

```

Function Union( $I, \ell$ )
   $i_0 := \operatorname{argmax}\{|P_i| : i \in I\};$ 
   $S := P_{i_0};$ 
  foreach  $i \in I \setminus \{i_0\}$  do
    foreach  $x \in P_i$  do  $tree[x] := S;$ 
     $S.MultiInsert(P_i);$ 
   $C := \emptyset;$ 
  foreach  $i \in I \setminus \{i_0\}$  do
    foreach  $(b, a) \in S.MultiPred(P_i)$  do
      if  $next[a] \neq b$  then  $C := C \cup \{(a, b)\};$ 
    foreach  $(a, b) \in S.MultiSucc(P_i)$  do
      if  $next[a] \neq b$  then  $C := C \cup \{(a, b)\};$ 
     $id^{-1}[i] := \text{null};$ 
   $id^{-1}[i_0] := \text{null};$ 
   $id[S] := \ell; id^{-1}[\ell] := S;$ 
  foreach  $(x, y) \in C$  do  $next[x] := y$  return  $C;$ 

```

Claim 3. Any sequence of *Union* operations takes $\mathcal{O}(n \log n)$ time in total.

Proof. Let us introduce a potential function $\Phi(\mathcal{P}) = \sum_{P \in \mathcal{P}} |P| \log |P|$. We shall prove that the running time of a single *Union* operation is proportional to the increase in potential. Clearly

$$0 \leq \Phi(\mathcal{P}) = \sum_{P \in \mathcal{P}} |P| \log |P| \leq \sum_{P \in \mathcal{P}} |P| \log n = n \log n,$$

so this suffices to obtain the desired $\mathcal{O}(n \log n)$ bound.

Let us consider a *Union* operation that merges partition classes of sizes $p_1 \geq p_2 \geq \dots \geq p_k$ to a single class of size $p = \sum_{i=1}^k p_i$. The most time-consuming steps of the algorithm are the operations on height-balanced trees, which, for single i , run in $O\left(\max\left(p_i, p_i \log \frac{p}{p_i}\right)\right)$ time. These operations are not performed for the largest set and for the remaining ones we have $p_i < \frac{1}{2}p$ (i.e. $\log \frac{p}{p_i} \geq 1$). This lets us bound the time complexity of the *Union* operation as follows:

$$\sum_{i=2}^k \max\left(p_i, p_i \log \frac{p}{p_i}\right) = \sum_{i=2}^k p_i \log \frac{p}{p_i} \leq \sum_{i=1}^k p_i \log \frac{p}{p_i} = \sum_{i=1}^k p_i (\log p - \log p_i) = p \log p - \sum_{i=1}^k p_i \log p_i,$$

which is equal to the increase in potential. □

6 By-Products of Cover Suffix Tree

In this section we present two additional applications of the Cover Suffix Tree. We show that, given $CST(w)$ (or CST of a word that can be obtained from w in a simple manner), one can compute in linear time all distinct primitively rooted squares in w and a linear representation of all the seeds of w , in particular, the shortest seeds of w . This shows that constructing this data structure is *at least as hard* as computing all primitively rooted squares and seeds. While there are linear-time algorithms for these problems [16, 20, 10] and [19], they are all complex and rely on the combinatorial properties specific to the repetitive structures they seek for.

Theorem 4. Assume that the Cover Suffix Tree of a word of length n can be computed in $T(n)$ time. Then all distinct primitively rooted squares in a word w of length n can be computed in $T(2n)$ time.

Proof. Let $0 \notin \Sigma$ be a special symbol. Let $\varphi : \Sigma^* \rightarrow (\Sigma \cup \{0\})^*$ be a morphism such that $\varphi(c) = 0c$ for any $c \in \Sigma$. We consider the word $w' = \varphi(w)0$, that is, the word w with 0-characters inserted at all its inter-positions, e.g. if $w = \mathbf{aabab}$ then $w' = 0\mathbf{a}0\mathbf{a}0\mathbf{b}0\mathbf{a}0\mathbf{b}0$.

Let us consider the set of explicit non-branching nodes of $CST(w')$ and select among them the nodes corresponding to even-length factors of w' starting with the symbol 0. It suffices to note that there is a one-to-one correspondence between these nodes and the halves of primitively rooted squares in w . \square




Figure 6: Seed of string **aaabaabaabaabaaba**.

Recall that a word u is a seed of w if u is a factor of w and w is a factor of some word y covered by u , see Fig. 6. The following lemma states that the set of all seeds of w has a representation of $\mathcal{O}(n)$ size, where $n = |w|$. This representation enables, e.g., simple computation of all shortest seeds of the word. By a range on a edge of a suffix tree we mean a number of consecutive nodes on this edge (obviously at most one of these nodes is explicit). Let w^R denote the reverse of the word w .

Lemma 6 ([18, 19]). *The set of all seeds of w can be split into two disjoint classes. The seeds from one class form a single (possibly empty) range on each edge of the suffix tree of w , while the seeds from the other class form a range on each edge of the suffix tree of w^R .*

We will show that given $CST(w)$ and $CST(w^R)$ we can compute the representation of all seeds from Lemma 6 in $\mathcal{O}(n)$ time. Let us recall auxiliary notions of quasiseed and quasigap, see [19].

By $first(u)$ and $last(u)$ let us denote $\min Occ(u)$ and $\max Occ(u)$, respectively. We say that u is a *complete cover* in w if u is a cover of the word $w[first(u), last(u) + |u| - 1]$. The word u is called a *quasiseed* of w if u is a complete cover in w , $first(u) < |u|$ and $n + 1 - last(u) < 2|u|$. Alternatively, w can be decomposed into $w = xyz$, where $|x|, |z| < |u|$ and u is a cover of y .

All quasiseeds of w lying on the same edge of the suffix tree with lower explicit endpoint v form a range with the lower explicit end of the range located at v . The length of the upper end of the range is denoted as $quasigap(v)$. If the range is empty, we set $quasigap(v) = \infty$. Thus a representation of all quasiseeds of a given word can be provided using only the quasigaps of explicit nodes in the suffix tree. It is known that computation of quasiseeds is the hardest part of an algorithm computing seeds:

Lemma 7 ([18, 19]). *Assume quasigaps of all explicit nodes of suffix trees of w and w^R are known. Then a representation of all seeds of w from Lemma 6 can be found in $\mathcal{O}(n)$ time.*

It turns out that the auxiliary data in $CST(w)$ and $CST(w^R)$ enable constant-time computation of quasigaps of explicit nodes. By Lemma 7 this yields an $\mathcal{O}(n)$ time algorithm for computing a representation of all the seeds of w . This is stated formally in the following theorem.

Theorem 5. *Assume that the Cover Suffix Tree of a word of length n can be computed in $T(n)$ time. Given a word w of length n , one can compute a representation of all seeds of w from Lemma 6 in $T(n)$ time. In particular, all the shortest seeds of w can be computed within the same time complexity.*

Proof. We show how to compute quasigaps for all explicit nodes of $CST(w)$. The computation for $CST(w^R)$ is symmetric. Note that $CST(w)$ may contain more explicit nodes than the suffix tree of the word. In this case, the results from any maximal sequence of edges connected by non-branching explicit nodes in $CST(w)$ need to be merged into a single range on the corresponding edge of the suffix tree.

By the definition of $cv(v)$, an explicit node v of $CST(w)$ is a complete cover in w if the following condition holds:

$$cv(v) = last(v) - first(v) + |v|.$$

Thus for checking whether an explicit node v of $CST(w)$ is a quasisseed of w it suffices to check whether this condition and the following equalities hold:

$$first(v) < |v|, \quad n + 1 - last(v) < 2|v|.$$

If v is not a quasisseed of w , we have $quasigap(v) = \infty$, otherwise we can assume that $quasigap(v) \leq |v|$.

Example 3. Consider the word w from Fig. 2, $n = 15$. The word **cacc** corresponds to an explicit node of $CST(w)$; we denote it by v . We have $cv(v) = 11$, $first(v) = 4$, $last(v) = 11$, and $last(v) - first(v) + |v| = 11$. Therefore **cacc** is a quasisseed of w , see also Fig. 1.

By Lemma 1, the condition for any node on the edge ending at v to be a complete cover in w is very simple:

$$\Delta(v) = 1.$$

Assume this condition is satisfied and consider any implicit node v' on this edge. Then v' is a quasisseed if both inequalities:

$$first(v) < k \quad \text{and} \quad n + 1 - last(v) < 2k$$

are satisfied. Thus in this case

$$quasigap(v) = \max(first(v) + 1, \lceil (n - last(v) + 2)/2 \rceil, |parent(v)| + 1).$$

Example 4. Consider the word w from Fig. 2. The word **cccacc** corresponds to an explicit node of $CST(w)$; we denote it by v . We have $cv(v) = 10$, $first(v) = 2$, $last(v) = 6$, and $last(v) - first(v) + |v| = 10$. Therefore **cccacc** is a quasisseed of w . Since $\Delta(v) = 1$, $quasigap(v)$ could be smaller than 6. However, $\lceil (n - last(v) + 2)/2 \rceil = 6$ and the above formula yields $quasigap(v) = 6$.

This concludes a complete set of rules for computing $quasigap(v)$ for explicit nodes of $CST(w)$. □

7 Conclusions

We have presented an algorithm which constructs a data structure, called the *Cover Suffix Tree*, in $\mathcal{O}(n \log n)$ time and $\mathcal{O}(n)$ space. The Cover Suffix Tree has been developed in order to solve the PARTIALCOVERS and ALLPARTIALCOVERS problem in $\mathcal{O}(n)$ and $\mathcal{O}(n \log n)$ time, respectively, but it also gives a well-structured description of the cover indices of all factors. Consequently, various questions related to partial covers can be answered efficiently. For example, with the Cover Suffix Tree one can solve in linear time a problem inverse to PARTIALCOVERS: find a factor of length between l and r that maximizes the number of positions covered. Also a similar problem to ALLPARTIALCOVERS problem, to compute for all lengths $l = 1, \dots, n$ the maximum number of positions covered by a factor of length l , can be solved in $\mathcal{O}(n \log n)$ time. This solution was actually given implicitly in the proof of Lemma 3.

An interesting open problem is to reduce the construction time to $\mathcal{O}(n)$. This could be difficult, though, since by the results of Section 6 this would yield alternative linear-time algorithms finding primitively rooted squares and computing seeds. The only known linear-time algorithms for these problems (see [16, 10, 11] and [19]) are rather complex.

References

- [1] A. Apostolico and A. Ehrenfeucht. Efficient detection of quasiperiodicities in strings. *Theor. Comput. Sci.*, 119(2):247–265, 1993.
- [2] A. Apostolico, M. Farach, and C. S. Iliopoulos. Optimal superprimitivity testing for strings. *Inf. Process. Lett.*, 39(1):17–20, 1991.
- [3] A. Apostolico and F. P. Preparata. Data structures and algorithms for the string statistics problem. *Algorithmica*, 15(5):481–494, 1996.
- [4] D. Breslauer. An on-line string superprimitivity test. *Inf. Process. Lett.*, 44(6):345–347, 1992.
- [5] G. S. Brodal, R. B. Lyngsø, A. Östlin, and C. N. S. Pedersen. Solving the string statistics problem in time $O(n \log n)$. In P. Widmayer, F. T. Ruiz, R. M. Bueno, M. Hennessy, S. Eidenbenz, and R. Conejo, editors, *ICALP*, volume 2380 of *Lecture Notes in Computer Science*, pages 728–739. Springer, 2002.
- [6] G. S. Brodal and C. N. S. Pedersen. Finding maximal quasiperiodicities in strings. In R. Giancarlo and D. Sankoff, editors, *CPM*, volume 1848 of *Lecture Notes in Computer Science*, pages 397–411. Springer, 2000.
- [7] M. R. Brown and R. E. Tarjan. A fast merging algorithm. *J. ACM*, 26(2):211–226, 1979.
- [8] M. Crochemore, C. Hancart, and T. Lecroq. *Algorithms on Strings*. Cambridge University Press, 2007.
- [9] M. Crochemore, L. Ilie, and W. Rytter. Repetitions in strings: Algorithms and combinatorics. *Theor. Comput. Sci.*, 410(50):5227–5235, 2009.
- [10] M. Crochemore, C. Iliopoulos, M. Kubica, J. Radoszewski, W. Rytter, and T. Waleń. Extracting powers and periods in a word from its runs structure. *Theoretical Computer Science*, doi: 10.1016/j.tcs.2013.11.018, 2013.
- [11] M. Crochemore, C. S. Iliopoulos, M. Kubica, J. Radoszewski, W. Rytter, and T. Waleń. Extracting powers and periods in a string from its runs structure. In E. Chávez and S. Lonardi, editors, *SPIRE*, volume 6393 of *Lecture Notes in Computer Science*, pages 258–269. Springer, 2010.
- [12] M. Crochemore and W. Rytter. *Jewels of Stringology*. World Scientific, 2003.
- [13] M. Farach. Optimal suffix tree construction with large alphabets. In *FOCS*, pages 137–143, 1997.
- [14] T. Flouri, C. S. Iliopoulos, T. Kociumaka, S. P. Pissis, S. J. Puglisi, W. Smyth, and W. Tyczyński. Enhanced string covering. *Theoretical Computer Science*, 506(0):102 – 114, 2013.
- [15] A. S. Fraenkel and J. Simpson. How many squares can a string contain? *J. Comb. Theory, Ser. A*, 82(1):112–120, 1998.
- [16] D. Gusfield and J. Stoye. Linear time algorithms for finding and representing all the tandem repeats in a string. *J. Comput. Syst. Sci.*, 69(4):525–546, 2004.
- [17] J. Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Inf. Process. Lett.*, 33(4):169–174, 1989.
- [18] C. S. Iliopoulos, D. W. G. Moore, and K. Park. Covering a string. *Algorithmica*, 16(3):288–297, 1996.
- [19] T. Kociumaka, M. Kubica, J. Radoszewski, W. Rytter, and T. Waleń. A linear time algorithm for seeds computation. In Y. Rabani, editor, *SODA*, pages 1095–1112. SIAM, 2012.
- [20] R. M. Kolpakov and G. Kucherov. Finding maximal repetitions in a word in linear time. In *FOCS*, pages 596–604. IEEE Computer Society, 1999.

- [21] G. Kucherov, Y. Nekrich, and T. A. Starikovskaya. Cross-document pattern matching. In J. Kärkkäinen and J. Stoye, editors, *CPM*, volume 7354 of *Lecture Notes in Computer Science*, pages 196–207. Springer, 2012.
- [22] Y. Li and W. F. Smyth. Computing the cover array in linear time. *Algorithmica*, 32(1):95–106, 2002.
- [23] D. Moore and W. F. Smyth. An optimal algorithm to compute all the covers of a string. *Inf. Process. Lett.*, 50(5):239–246, 1994.
- [24] J. S. Sim, K. Park, S. Kim, and J. Lee. Finding approximate covers of strings. *Journal of Korea Information Science Society*, 29(1):16–21, 2002.
- [25] E. Ukkonen. On-line construction of suffix trees. *Algorithmica*, 14(3):249–260, 1995.